



Requesting a Project Columbia Allocation

Jim Fischer/GSFC/930

**Presented to the High Performance Computing Workshop
Building 26, GSFC**

November 22, 2004





Requesting a Project Columbia Allocation

Jim Fischer/GSFC/930

**Presented to the High Performance Computing Workshop
Building 26, GSFC**

November 22, 2004



Project Columbia History

- ❑ Idea started in August 2003 as a collaboration between the Earth Science Enterprise and Aeronautics Technology Enterprise
- ❑ The first SGI Altix system 'Kalpana' arrived in April 2004
- ❑ Initiated partnership discussions with Intel and SGI in April 2004
- ❑ Proposed by ARC to NASA Executive Committee mid-May 2004
- ❑ Approved by OMB and Congress in June 2004
- ❑ Funded in June 2004 by NASA
- ❑ Delivery and installation from July to mid-October
- ❑ Dedicated October 26, 2004
- ❑ Now ranks #2 on Top500 list sustaining 51.8 Teraflops on LinPack, compared to the, now #3, Earth Simulator's 34.8 Teraflops.

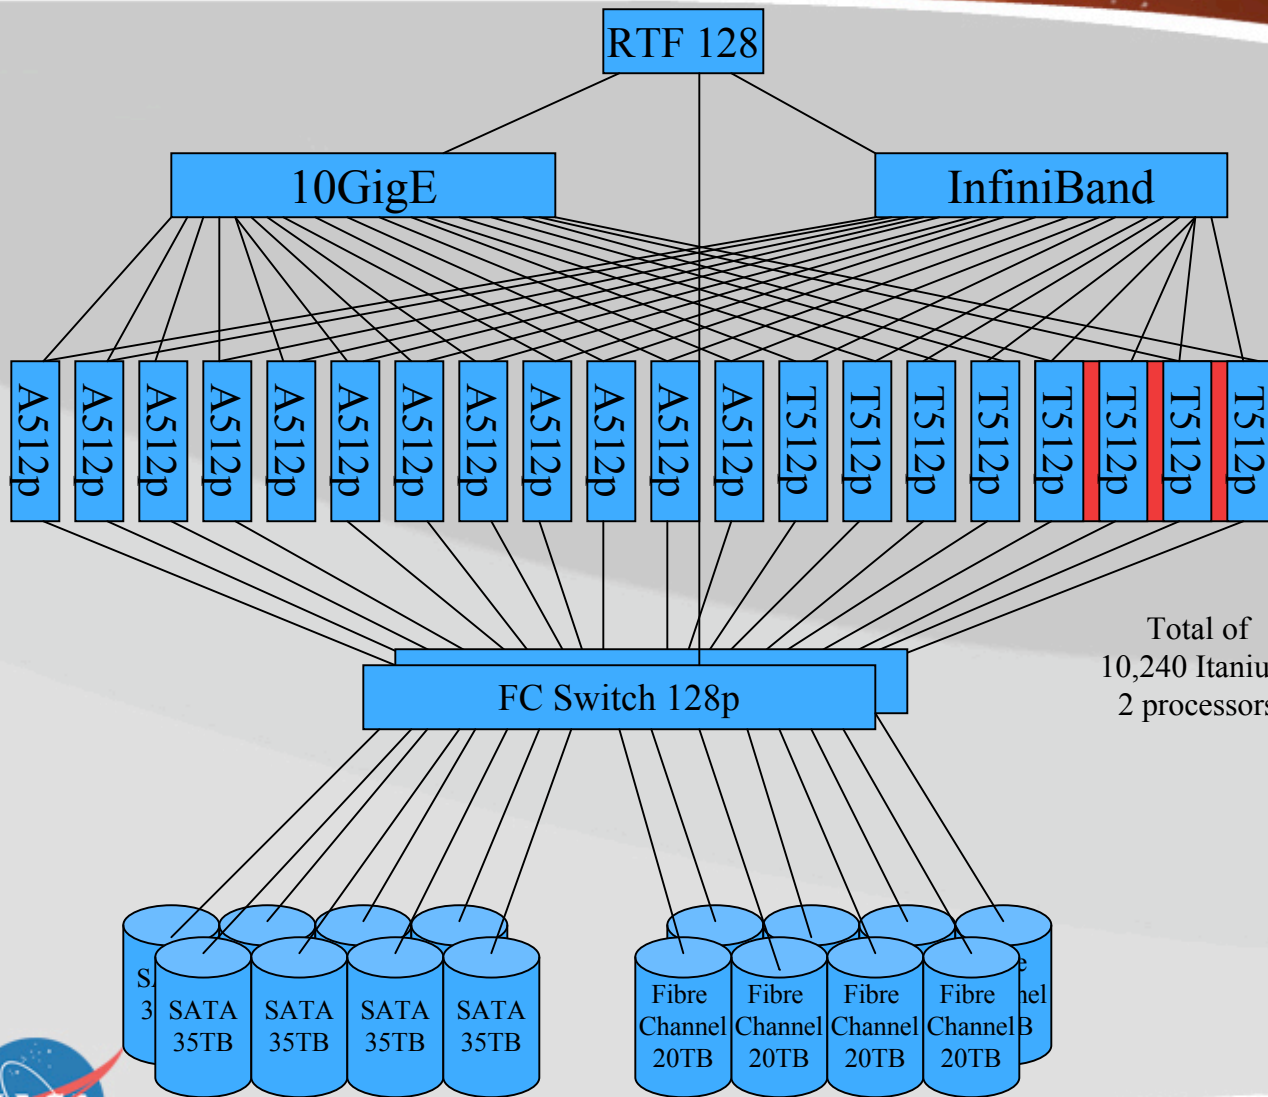


Project Columbia Characteristics

- ☐ Partnership between NASA and industry to significantly enhance the national compute capability
- ☐ Unique opportunity that met both industry and NASA objectives
- ☐ In line with HECRTF findings
- ☐ National asset available to multiple agencies through competitive selection process
- ☐ Unique capability built from proven technologies
- ☐ Asset enhances capability that was not currently being targeted by other leadership class systems
- ☐ Provides NASA with increased capabilities
- ☐ Includes networking component to connect NASA centers for scientific research



Components



Total of
10,240 Itanium
2 processors

Front End

- 128p Altix 3700 (RTF)

Networking

- InfiniBand Switch (288port)
- InfiniBand (4 per 3700, 8 per 3700Bx2))
- 10GigE Switch 32-port
- 10GigE Cards (1 Per 512p)
- GigE Switch 384-port, 96-port
- GigE (12 per 512)
- Altix 3700Bx2 2048 Numalink Kit

Compute Nodes

- Altix 3700 12x512p
- Altix 3700Bx2 8x512p

Storage Area Network

- Brocade Switch 2x128port

Storage (440 TB)

- FC RAID 8x20 TB (8 Racks)
- SATARAID 8x35TB (8 Racks)

Mass Storage (17 PB)

- 17 PB MSS



Internetworking

- ❑ "NUMAlink" cables packed with copper wires connect groups of 16 cabinets into single 512-processor systems.
- ❑ Twenty of these 512 processor systems, connected with InfiniBand networking equipment from Voltaire, make up Columbia.


(Photo credit: Stephen Shankland)



Cooling System

- ❑ The forthcoming BX2 model of SGI's Altix 3700 packs 64 Itanium 2 processors into one cabinet, compared to 32 for its predecessor.
- ❑ The resulting heat led SGI to offer the system with liquid cooling built into the back door.

(Photo credit: Stephen Shankland)




Disk Storage System


- Columbia's disk storage system, built by SGI, holds 440 terabytes of data.

*(Photo credit:
Stephen Shankland)*





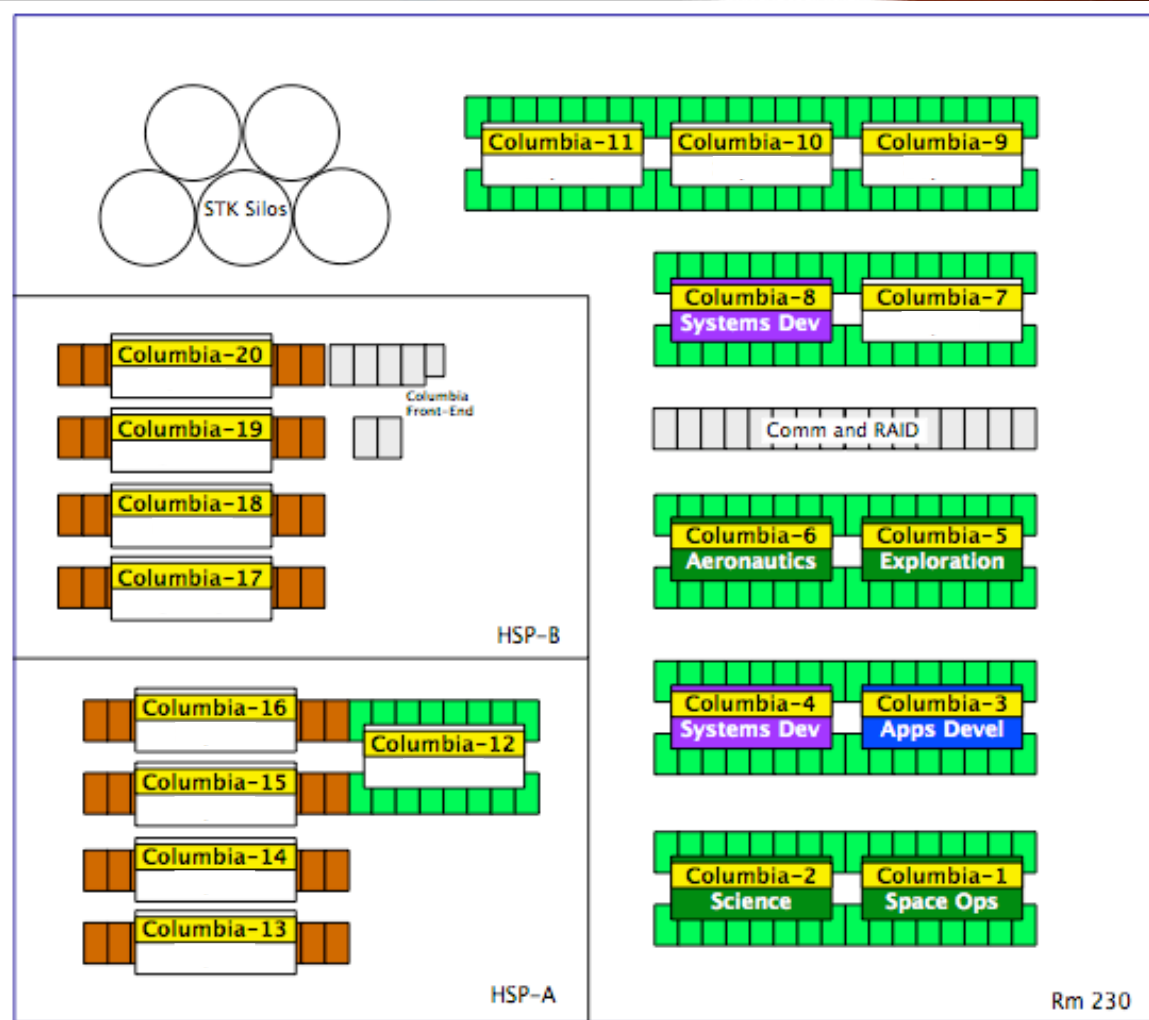
Columbia Footprint



Located in the NAS Building at NASA/ARC



Allocation Approach circa 9/04



Network Implementation Plan

	Oct 04	Jul 05	Jan 06	Jul 06
	Phase 0	Phase 1	Phase 2	Phase 3
	1 Gbps	10 Gbps	10 Gbps	10 Gbps
ARC	x	x		
JPL	x	x		
GSFC		x		
Starlight		x		
MSFC			x	
LRC			x	
GRC			x	
JSC			x	
NGIX-West		x		
NGIX-East		x		
KSC				x
WSTF				x




Opportunity and Challenge

- ☐ Columbia is funded by all four NASA Mission Directorates
- ☐ Funding is through FY07
- ☐ No funding beyond FY07 in the current plan
- ☐ Agency is pushing a “service pool” model for the future funding
- ☐ 28% of Columbia allocated to SMD
- ☐ Earth-Sun modeling has been the prime usage of the system
- ☐ Demand exceeds supply
- ☐ Science Mission needs to productively use the system while articulating and advocating the “need”




Earth-Sun Applications on Columbia

Space Weather Modeling Framework




1/4° Finite Volume General Circulation Model

NASA fvGCM
Clouds [fraction] : Total Precipitable Water [kg/m²]
2004 SEP 05 21:00Z



Estimating the Circulation and Climate of the Ocean (ECCO)




Project Columbia Governance

- ❑ NASA Deputy Administrator has directed formation of the Project Columbia “Board of Directors (BOD),” organized by CIO, with representation from all four Mission Directorates
- ❑ It has been established and is chaired by the Deputy CIO
- ❑ It serves as a forum for discussions on high level allocation, budget advocacy, and Interagency relationship/partnership
- ❑ The Science Mission Directorate is represented by Joe Bredekamp
- ❑ The BOD has given the Science Mission Directorate the leadership role to develop OneNASA calls for proposals for Project Columbia “Guest Investigators,” as well as calls to give other agencies opportunities to request Columbia time
- ❑ Calls such as the MAP NRA contain mechanisms that allow requests for Columbia time



Calls for Proposals

- A “**Dear Colleague**” letter will be released every 6 months via Mission Directorate email distribution lists (Exploration, Science, Aeronautics, Space Operations) Announcing the opportunity to request Columbia resources.
- NOT an NRA - no \$s or FTEs awarded, only Columbia resources
- An Allocation Panel will receive proposals, oversee a rating and ranking process, and develop recommendations
- The Columbia Board of Directors will finalize the allocations, which are good for a year.



Annual Allocation to Science

- The first call for Guest Investigator proposals is poised for release
- Only Science Mission Directorate time will be offered in the first call
- The amount of time to be offered has not been finalized

Project Columbia Cycles **Per Year**

Total CPU-hours per year		89702400
Reserved for System Testing	10%	(8,970,240)
Reserved for Interagency	5%	(4,485,120)
Reserved for Education	2%	(1,794,048)
Reserved for system inefficiency	10%	(8,970,240)
Reserved for 2048 (1024 for now)	10%	(8,970,240)
Available for Missions	63%	56,512,512

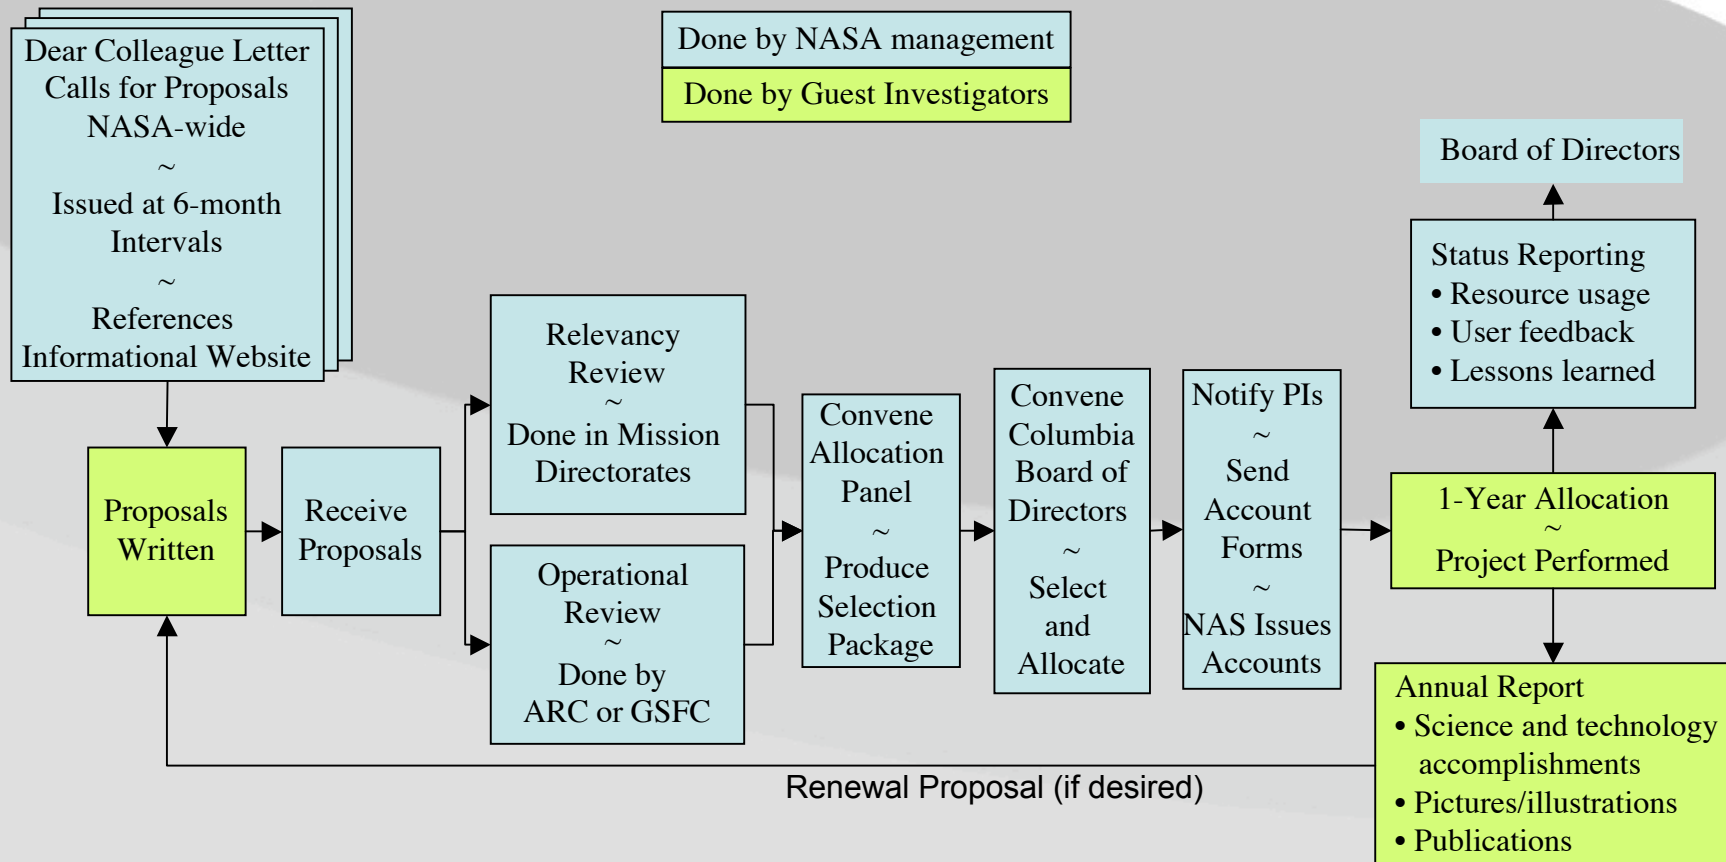
Available for Missions **56,512,512**

Space Ops & NESC	14.4%	8,120,674
Aeronautic	11.3%	6,385,803
Exploration	46.2%	26,133,807
Science	28.1%	15,872,227
Total	100.0%	56,512,512

Estimated Price per CPU hour **0.90**



Guest Investigator Proposal Process



Summary

- ❑ Project Columbia is an opportunity to organize a more coherent OneNASA strategy for high-end computing
- ❑ It will be around for quite a while
- ❑ Its appearance is disrupting the GSFC high-end computing status quo
- ❑ A new strategic relationship is needed among NASA's high-end computing assets for the benefit of science
- ❑ Columbia's use may expose opportunities for complementary technology programs
- ❑ The first OneNASA call for proposals for Project Columbia resources is about to be released



Computing as an Instrument

